arXiv:1604.02614v1 [math.NA] 9 Apr 2016

On the performance of exponential integrators for problems in magnetohydrodynamics

Lukas Einkemmer[a,*], Mayya Tokman[b], John Loffeld[c]

[a] *University of Innsbruck, Austria*
[b] *University of California, Merced, United States of America*
[c] *Lawrence Livermore National Laboratory, Livermore, United States of America*

Abstract

Exponential integrators have been introduced as an efficient alternative to explicit and implicit methods for integrating large stiff systems of differential equations. Over the past decades these methods have been studied theoretically and their performance was evaluated using a range of test problems. While the results of these investigations showed that exponential integrators can provide significant computational savings, the research on validating this hypothesis for large scale systems and understanding what classes of problems can particularly benefit from the use of the new techniques is in its initial stages. Resistive magnetohydrodynamic (MHD) modeling is widely used in studying large scale behavior of laboratory and astrophysical plasmas. In many problems numerical solution of MHD equations is a challenging task due to the temporal stiffness of this system in the parameter regimes of interest. In this paper we evaluate the performance of exponential integrators on large MHD problems and compare them to a state-of-the-art implicit time integrator. Both the variable and constant time step exponential methods of EpiRK-type are used to simulate magnetic reconnection and the Kevin-Helmholtz instability in plasma. Performance of these methods, which are part of the EPIC software package, is compared to the variable time step variable order BDF scheme included in the CVODE (part of SUNDIALS) library. We study performance of the methods on parallel architectures and with respect to magnitudes of important parameters such as Reynolds, Lundquist, and Prandtl numbers. We find that the exponential integrators provide superior or equal performance in most circumstances and conclude that further development of exponential methods for MHD problems is warranted and can lead to significant computational advantages for large scale stiff systems of differential equations

such as MHD.

1. Introduction

The focus of this article lies at the intersection of computational plasma physics and numerical analysis. The system of resistive magnetohydrodynamics (MHD) equations is used extensively in both laboratory and astrophysical plasma physics [16]. MHD simulations are employed to study large scale plasma dynamics in applications ranging from design of the next generation of fusion devices to modeling astrophysical jets (e.g. [26, 10, 53]). Since experiments in this field usually impose a high cost and the ability to conduct observations can be very limited, computational modeling in general, and numerical MHD in particular, have become very important components of research in plasma physics. While MHD theory is widely used to study plasma dynamics, the complexity of the MHD equations makes their numerical solution a challenging task. In particular, the MHD system describes a wide range of temporal scales. In applications such as large scale modeling of tokamaks or modeling of solar eruptions, the physical processes of interest evolve on the time scale of the shear Alfvén waves, which is much slower than the fastest mode in the system - the fast magnetosonic wave - and much faster than the slow magnetosonic wave and the resistive processes. Such temporal stiffness coupled with the complexity of the equations poses a double-edged challenge to developing a time integrator for MHD. From one perspective, the stiffness implies that explicit time integrators will suffer from prohibitively severe restrictions on the time step size due to stability considerations. On the other hand the inherent three-dimensional nature of the processes MHD describes and the complexity of spatial operators describing the dynamics makes evaluation of forcing terms computationally expensive. The latter difficulty served as one of the major reasons most large scale MHD codes came to rely extensively on explicit time integrators even for problems where resistivity plays a major role. However, the temporal stiffness of the equations forced researchers working on these software packages to turn to at least some form of implicitness to overcome the small time steps sizes of explicit integrators [22]. Production codes used

*Corresponding author
Email addresses: lukas.einkemmer@uibk.ac.at (Lukas Einkemmer), mtokman@ucmerced.edu (Mayya Tokman), loffeld1@llnl.gov (John Loffeld)

in fusion research and astrophysics [44, 4, 15] started embedding some form of implicitness into their time integration schemes. The NIMROD code, for instance, employs an implicitly integrated term in the momentum equation derived by linearizing equations around a static magnetic field and using it to stabilize stiff wave components [45, 44]. In addition to this stabilizing term the overall time integration strategy also uses operator splitting and a staggered temporal grid. The choice of the implicit terms in such semi-implicit schemes is also influenced by the availability of efficient algebraic solvers used to approximate solutions to the associated linear systems. The complexity of the resulting overall time integration method makes it very difficult to use rigorous analysis to study stability and convergence. Such an approach also precludes or makes it challenging to derive error estimators necessary for employing adaptive time stepping strategies or using the time integration as part of the uncertainty quantification methods. Other partially implicit approaches such as alternating direction implicit (ADI) methods [27, 42] or implicit-explicit techniques [24] are also limited as they try to maintain a balance between stabilizing the time integration while maintaining the accuracy of the solution [7]

The difficulties associated with partially implicit methods and the ever pressing need for more efficient and accurate algorithms for high-performance computing platforms gave rise to the efforts to develop fully implicit time integrators for resistive MHD [38, 37, 7, 30]. A typical approach to constructing a fully implicit integrator involves using an implicit scheme such as, for instance, BDF and coupling it with a Newton-Krylov iteration to solve the implicit equations within each time step. In [38, 37], for example, the resistive MHD equations are solved using a variable order variable time step BDF method from the CVODE portion [8] of the SUNDIALS package [17]. While implicit methods possess better stability properties compared to explicit techniques, they are also affected by the stiffness of the equations. In an implicit integrator temporal stiffness manifests itself in the stiff linear systems that have to be solved as part of each Newton iteration. A key to obtaining an efficient fully implicit Newton-Krylov solver is in developing effective scalable preconditioners for these stiff linear systems. For a complex anisotropic system such as MHD construction of a sufficiently efficient preconditioner is simultaneously a central and a challenging task that requires a significant investment of time and effort. If the parameter regime is changed or an extra term is added to the system, it is possible that a new preconditioner has to be developed. Often the structure of the equations

is exploited to build a physics based preconditioner [7] which makes the overall integrator highly dependent on the equations, thus precluding one from using off-the-shelf time integration software packages.

In recent years exponential integrators have emerged as a promising alternative to explicit and implicit methods for solving evolution equations. A significant body of research has been accumulated in which these methods are investigated both from a theoretical (see e.g. [21] for a review article) as well as from a numerical point of view (see e.g. [51] and [28]). This research includes the construction of exponential integrators that have been tailored to a given differential equation (see e.g. [18]) as well as such integrators that can be applied to a class of problems (see e.g. [20, 47]). There is strong evidence to suggest that exponential integrators can offer computational advantages compared to standard integrators particularly for large scale complex problems where no efficient preconditioner is available. Exponential integration can also be used as part of an Implicit-Exponential (IMEXP) scheme [49] for problems where an efficient preconditioner has been constructed for some terms in the equations.

Exponential integrators were first shown to be a promising approach in the context of the MHD equations in [46, 50]. Since then a number of more efficient and sophisticated exponential time integrators have been developed. Exponential propagation iterative methods of Runge-Kutta type (EpiRK) were first introduced in [47] for general nonlinear systems of ordinary differential equations. The EpiRK framework [47, 51, 49] for constructing exponential time integrators has been designed to reduce computational complexity in each time step and provide a flexible ansatz that allows construction of particularly efficient exponential schemes [51, 48, 35, 49]. The most efficient EpiRK methods have been implemented as part of the Exponential Integrator Collection (EPIC) package, a C++/MPI software developed for serial and parallel computational platforms (available for download by request to mtokman@ucmerced.edu). Performance of EPIC integrators has been tested on a suit of numerical examples and it was shown that the EpiRK schemes can achieve superior or equal performance compared to the BDF integrator found in the CVODE library [28, 29].

However, up to now most numerical tests of exponential methods found in the literature considered problems of small to medium complexity (see, for example, [28], [25], or [12]) where in many instances efficient preconditioners can be constructed. The goal of this paper is to apply the EpiRK integrators found in the EPIC library to the problem of solving the resis-

tive magnetohydrodynamics (MHD) equations. General preconditioners for such systems have, for example, been investigated in [37] and it was found that constructing an efficient preconditioner for such system is challenging and the the performance gain from standard preconditioning techniques for several dimensional examples is rather modest under most circumstances. Clearly a problem specific preconditioner or a tailor-made time integrator for a specific system of equations will be the most efficient approach in the majority of cases, but developing such a scheme requires extensive effort and time. The goal of this paper is to investigate what computational savings can be achieved with a general exponential integrator compared to an unpreconditioned implicit scheme. Thus we will perform a comparison of EPIC integrators with the BDF methods from the CVODE library without a preconditioner. We aim to demonstrate that exponential integration is a viable technique to use in integrating MHD equations.

The paper is organized as follows. In section 2 we provide an introduction to exponential integrators in general as well as to the specific EpiRK methods we use in the numerical examples and their implementation in EPIC. In section 3 we describe the MHD equations which are solved numerically in section 4 for a reconnection problem, and in section 5 for the Kelvin-Helmholtz instability. Implementation details are discussed in section 6. Finally, our conclusions are presented in section 7.

2. Exponential integrators

Let us consider the following initial value problem

$$\begin{aligned} y'(t) &= F(y(t)) \\ y(0) &= y_0 \end{aligned} \tag{1}$$

which is assumed to be large and stiff. Before proceeding to describe the class of numerical methods we are interested in, let us use a Taylor expansion to write the right hand side of (1) in the following form

$$y'(t+h) = F(y(t)) + J(y(t))\left[y(t+h) - y(t)\right] + R(y(t+h)),$$

where $J(y(t)) = D_y F(y(t))$ is the Jacobian and the remainder function is given by

$$R(y(t+h)) = F(y(t+h)) - F(y(t)) - J(y(t))\left[y(t+h) - y(t)\right].$$

Now using the Gröbner-Alekseev formula (or an integrating factor e^{-hJ}), we obtain

$$y(t+h) = y(t) + \left(e^{hJ} - I\right)(hJ)^{-1} hF(y(t)) + \int_0^1 e^{hJ(1-s)} R(y(t+sh))\,\mathrm{d}s.$$

This integral representation of the solution can be used as a base for constructing an exponential integrator as follows. First, the nonlinear integral in (2) has to be approximated in a suitable manner. If this is done using a polynomial approximation of the function $R(y(t+sh))$ we express the integral as a linear combination of the entire functions defined by

$$\begin{aligned} \varphi_0(hJ) &= e^{hJ} \\ \varphi_k(hJ) &= \int_0^1 e^{(1-\theta)hJ} \frac{\theta^{k-1}}{(k-1)!}\,\mathrm{d}\theta \quad \text{for } k \geq 1. \end{aligned} \tag{2}$$

The numerical solution is then approximated as a linear combination of the products of matrix functions and vectors $\varphi_k(chJ)v$ where c is a constant. These approximations constitute the main computational cost of an exponential method, therefore choosing an appropriate algorithm is essential to obtaining an efficient integrator. A number of methods have been proposed in the literature to approximate either a general or a specific (e.g. rational or exponential) function of a matrix or its product with a given vector (e.g. [9, 34, 1, 23, 6] or see the review in [21]). However, many of these techniques are applicable only to small matrices J [32, 43]. In this paper we are concerned with large scale problems such as MHD, where the dimensionality of the system (1) is high (e.g. number of unknowns in MHD equations multiplied by the number of grid points in two- or three-dimensions), and thus standard techniques for approximating a product of a matrix function φ_k with a vector are either computationally infeasible or prohibitively expensive. One possibility is to use a polynomial approximation of φ_k. To that end we have to chose m interpolation nodes on some compact set $K \subset \mathbb{C}$. Chebyshev nodes are an obvious choice, however, they suffer from the disadvantage that to compute the interpolation for $m+1$ points we have to reevaluate all the matrix-vector products already computed. To remedy this, Leja points have been proposed which share many of the favorable properties of the Chebyshev nodes but can be generated in sequence [6, 31, 5, 3]. Another advantage of interpolation at Leja points is its modest memory requirements. This makes it an attractive alternative

for computer systems where memory is limited (such as graphic processing units, see [12]). The main disadvantage of the method, however, is that (at least) some approximation of the spectrum of J has to be available. Since approximating the spectrum of a Jacobian matrix in MHD can be a very computationally expensive task, in this paper we focus exclusively on a more general approach based on the Krylov projection algorithm [2, 9].

Krylov methods have been extensively used to approximate the inverse of a matrix [41] and in 1987 were employed by Van der Vorst to approximate a general function of a matrix and a vector product [9]. The algorithm works by iteratively computing a projection of $\varphi_k(chJ)v$ onto a Krylov subspace

$$K_m(chJ, v) = \text{span}\left\{v, (chJ)v, \ldots, (chJ)^{m-1}v\right\}$$

using the Arnoldi algorithm [2]. The dimension of the Krylov subspace m is determined dynamically by estimating the residuals in the course of the iteration [40, 19]. For each m a matrix V_m is formed that consists of the orthonormal vectors v_i $(i = 1, ..., m)$ that form the basis of $K_m(chJ, v)$. A side product of the Arnoldi iteration is the $m \times m$ matrix $H_m = chV_m^T JV_m$ which represents the projection of chJ onto the Krylov subspace. A general product of the matrix function f with a vector is then approximated as

$$f(chJ)b \approx V_m \left(V_m^T \varphi_k(chJ)V_m\right) V_m^T v \approx V_m \varphi_k(H_m)V_m^T v.$$

Note that H_m is a matrix of dimension $m \times m$ and therefore the application of $\varphi_k(H_m)$ to a vector can be easily computed by using any of the standard methods (such as Pade approximation or methods based on polynomial interpolation [32, 21]).

Since evaluations of $\varphi_k(chJ)v$ constitute the most computationally expensive part of the exponential integrator, special care has to be taken to ensure that each time step requires as few of these evaluations as possible and the most efficient algorithm is chosen for these approximations. The EpiRK framework [47, 48] is designed to construct integrators with a reduced number of $\varphi_k(chJ)v$ evaluations and with φ_k functions chosen to speed up the Krylov iteration. A particularly efficient class of EpiRK methods [52] also takes advantage of an adaptive version of the Krylov iteration introduced in [34]. The adaptive Krylov based EpiRK schemes link the construction of a time integrator with the choice of the method to approximate $\varphi_k(chJ)v$ to build a more efficient overall time stepping scheme in the following way.

The computational cost of the Krylov algorithm to approximate terms of type $\varphi_k(chJ)v$ scales as $O(m^2)$ with the size of the Krylov basis m. Clearly, m depends on the eigenvalues of J, vector v as well as the values of c and h. As c or h is increased, so is the size of the basis. The EpiRK framework allows one to reduce the coefficients c and choose the better suited functions φ_k in the derivation of a time integrator. The adaptive Krylov technique allows replacing each computation of $\varphi_k(chJ)v$ with several evaluations of $\varphi_k(\tau_i chJ)v$ where $0 < \tau_i < 1$ for $i = 0, 1, 2, ...P$. Since the complexity of each of these evaluations now scales as $O(m_i^2)$, in most cases increased efficiency is observed as $O(m^2)$ can be less than $O(m_1^2)+O(m_2^2)+...+O(m_P^2)$.

A general EpiRK method can be written as

$$Y_i = y_n + a_{i1}\psi_{i1}(g_{i1}hJ_{i1})hF_n + \sum_{j=2}^{i-1} a_{ij}\psi_{ij}(g_{ij}hJ_{ij})h\Delta^{(j-1)}R(y_n), \quad i = 1, .., s-1$$

$$y_{n+1} = y_n + b_1\psi_{s1}(g_{s1}hJ_{s1})hF_n + \sum_{j=2}^{s} b_j\psi_{sj}(g_{sj}hJ_{sj})h\Delta^{(j-1)}R(y_n)$$

where the divided differences $\Delta^{(h-1)}R(y_n)$ are computed by using the nodes $y_n, Y_1, \ldots, Y_{s-1}$. Different choices of functions ψ_{ij} and matrices J_{ij} result in different classes of EpiRK methods such as unpartitioned, partitioned or hybrid exponential or implicit-exponential (IMEXP) integrators [49]. In this paper we employ the schemes with the choice of $J_{ij} = J_n$ and ψ_{ij} set as linear combinations of φ_k functions

$$\psi_{ij}(z) = \sum_{k=1}^{s} p_{ijk}\varphi_k(z).$$

As in the Runge-Kutta case, to construct an efficient integrator the coefficients a_{ij}, g_{ij}, b_j and p_{ijk} have to be chosen subject to the appropriate order conditions. In this paper, we use the variable time step fifth order method EpiRK5P1

$$\begin{aligned} Y_1 &= y_0 + a_{11}\varphi_1(g_{11}hA_0)hF(y_0) \\ Y_2 &= y_0 + a_{21}\varphi_1(g_{21}hA_0)hF(y_0) + a_{22}\varphi_1(g_{22}hA_0)hR(Y_1) \\ y_1 &= y_0 + b_1\varphi_1(g_{31}hA_0)hF(y_0) + b_2\varphi_1(g_{32}hA_0)hR(Y_1) + b_3\varphi_3(g_{33}hA_0)h(-2R(Y1) + R(Y2)), \end{aligned} \tag{3}$$

that has been derived in [51] (see Table 1 for the coefficients; this method is

also employed in the performance comparison given in [28]) and the constant time step fourth order EpiRK4 method proposed in [36] given by

$$
\begin{aligned}
Y_1 &= y_n + \tfrac{1}{2}\varphi_1(\tfrac{1}{2}h_nJ_n)h_nF(y_n) \\
Y_2 &= y_n + \tfrac{2}{3}\varphi_1(\tfrac{2}{3}h_nJ_n)h_nF(y_n) \\
y_{n+1} &= y_n + \varphi_1(h_nJ_n)h_nF(y_n) + \left(32\varphi_3(h_nJ_n) - 144\varphi_4(h_nJ_n)\right)h_nR(Y_1) \\
&\quad + \left(-\tfrac{27}{2}\varphi_3(h_nJ_n) + 81\varphi_4(h_nJ_n)\right)h_nR(Y_2)
\end{aligned}
\tag{4}
$$

Table 1: Coefficients of the fifth-order method EPIRK5P1.

$$
\begin{bmatrix} a_{11} & & \\ a_{21} & a_{22} & \\ b_1 & b_2 & b_3 \end{bmatrix} = \begin{bmatrix} 0.35129592695058193092 & & \\ 0.84405472011657126298 & 1.6905891609568963624 & \\ 1.0 & 1.2727127317356892397 & 2.2714599265422622275 \end{bmatrix}
$$

$$
\begin{bmatrix} g_{11} & & \\ g_{21} & g_{22} & \\ g_{31} & g_{32} & g_{33} \end{bmatrix} = \begin{bmatrix} 0.35129592695058193092 & & \\ 0.84405472011657126298 & 0.5 & \\ 1.0 & 0.71111095364366870359 & 0.62378111953371494809 \end{bmatrix}
$$

Both of these EpiRK schemes utilize the adaptive Krylov technique for the evaluations of $\varphi_k(chJ)v$ in order to improve performance. It is important to note that while both of these are three-stage methods, EpiRK4 requires only two Krylov projections to be executed per time step since the scale invariance of the Arnoldi iteration allows simultaneous calculation of the Krylov bases for Y_1 and Y_2 and one more Krylov projection to evaluate y_{n+1} (e.g. see [36] for details). EPIRK5P1 requires three Krylov projections to be performed each time step [52]. In addition EPIRK4 is a so-called stiffly accurate method while EPIRK5P1 is derived using classical order conditions. The advantages of the EPIRK5P1 method, however, include the higher order and the ability to design embedded methods for the automatic time step control that do not require an additional Krylov projection. The adaptive time step control version of the methods of type EPIRK4 is still under development.

3. The resistive magnetohydrodynamics equations

The most fundamental theoretical description of a classical plasma comes from the kinetic equation. This so called Vlasov equation (the collisionless case) or Boltzmann equation (when collisions are of physical significance)

describes the time evolution of a particle-density in the $3 + 3$ dimensional phase space (the first three dimensions correspond to the space dependence while the remaining correspond to the velocity dependence of the particle-density). While a number of simulations of different plasma phenomena have been conducted with this approach, due to the high dimensionality of the phase space, a lower-dimensional approximation is usually used to render such simulations feasible. A similar discussion holds true for the gyrokinetic approximation that can be employed in plasmas where a strong external magnetic field along a given axis is present. In this case the phase space of the Vlasov equation is reduced to $3 + 2$ dimensions by averaging over the gyro motion. This procedure yields a good approximation under the assumption of low frequency as compared to the cyclotron frequency. For a more detailed discussion of gyrokinetic models, see e.g. [14] or [13].

However, in many applications (such as magnetic confinement fusion, spheromak experiments, and astrophysical plasmas) the timescales of interest are sufficiently long and/or the full three dimensional model is necessary to describe the physical phenomena. In that case the kinetic approach is usually not feasible (even on modern day supercomputers) and thus further simplifications have to be introduced. For the magnetohydrodynamics (MHD) equations, which we will describe in the remainder of this section, the assumption is made that, to a good approximation, the distribution in the velocity direction is Maxwellian; in other words, it is assumed that each sufficiently small volume in the plasma is in thermodynamic equilibrium. For a general overview of kinetic and MHD models (including a derivation of the MHD equations from the Vlasov equation) see e.g. [33]. Numerical computations in the context of the MHD model discussed in this paper are performed, for example, in [38], [39], and [37].

If the assumption is made that the plasma considered is in thermodynamic equilibrium, the equations of motion in a three dimensional phase space (a so called fluid model) can be derived. These are given by (in dimensionless units)

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \boldsymbol{v}) = 0 \tag{5}$$

$$\rho \frac{\partial \boldsymbol{v}}{\partial t} + \rho \boldsymbol{v} \cdot \nabla \boldsymbol{v} = \boldsymbol{J} \times \boldsymbol{B} - \nabla p \tag{6}$$

which we refer to as the continuity and momentum equation respectively. The equations are written in terms of the density ρ, the fluid velocity $\boldsymbol{v}$,

the magnetic field $\boldsymbol{B}$, the electric current density $\boldsymbol{J}$, and the pressure p. The dynamics is determined by the Lorentz force (the $\boldsymbol{J} \times \boldsymbol{B}$ term) and the pressure gradient force (the ∇p term). These fluid equations have to be coupled to an appropriate model of the electric field. Note, however, that if the (ideal) Ohm's law is assumed to hold, i.e.

$$\boldsymbol{E} + \boldsymbol{v} \times \boldsymbol{B} = 0,$$

then the electric field $\boldsymbol{E}$ can be eliminated from Maxwell's equation. This leaves us with the following system

$$\begin{aligned} &\frac{\partial \boldsymbol{B}}{\partial t} = -\nabla \times (\boldsymbol{v} \times \boldsymbol{B}) \\ &\boldsymbol{J} = \nabla \times \boldsymbol{B} \\ &\nabla \cdot \boldsymbol{B} = 0. \end{aligned} \tag{7}$$

The first of these equations yields the time evolution of $\boldsymbol{B}$, the second can be used to eliminate $\boldsymbol{J}$ from equation (6), and the third is the familiar solenoidal constraint imposed on the magnetic field.

A commonly employed approach to close these equations (see e.g. [39]) is to supplement them with the following equation of state

$$e = \frac{p}{\gamma - 1} + \frac{\rho}{2} v^2 + B^2,$$

where the time evolution of the energy density e is given by

$$\frac{\partial e}{\partial t} = \nabla \cdot \left((e + p + \tfrac{1}{2} B^2) \boldsymbol{v} - \boldsymbol{B} (\boldsymbol{B} \cdot \boldsymbol{v}) \right). \tag{8}$$

Collectively, the equations (5)-(8) in the variables $(\rho, \boldsymbol{v}, \boldsymbol{B}, e)$ yield a first-order system of 8 differential equations in 8 variables and are called the ideal magnetohydrodynamics equations (or the ideal MHD equations).

For the purpose of performing the spatial discretization, these equations are often cast into the so-called divergence form. Then, the equations of motion read as (see e.g. [39])

$$\frac{\partial U}{\partial t} + \nabla \cdot F(U) = 0$$

with state vector

$$U = \begin{bmatrix} \rho \\ \boldsymbol{v} \\ \boldsymbol{B} \\ e \end{bmatrix}$$

and

$$F(U) = \begin{bmatrix} \rho \boldsymbol{v} \\ \rho \boldsymbol{v} \otimes \boldsymbol{v} + (p + \frac{1}{2}B^2) I - \boldsymbol{B} \otimes \boldsymbol{B} \\ \boldsymbol{v} \otimes \boldsymbol{B} - \boldsymbol{B} \otimes \boldsymbol{v} \\ (e + p + \frac{1}{2}B^2)\boldsymbol{v} - \boldsymbol{B}(\boldsymbol{B} \cdot \boldsymbol{v}) \end{bmatrix}, \tag{9}$$

where we have denoted the tensor product by using the $\otimes$ symbol.

In this paper, as in [38] and [39], we will consider a slightly more general class of equations which, in addition to the dynamics discussed so far, includes dissipative effects (due to particle collisions in the plasma). To that end the hyperbolic flux vector $F(U)$ is extended in [38] by a diffusive part given by

$$F_d(U) = \begin{bmatrix} 0 \\ Re^{-1}\boldsymbol{\tau} \\ S^{-1}\left(\eta \nabla B - \eta (\nabla B)^{\mathrm{T}}\right) \\ Re^{-1}\boldsymbol{\tau} \cdot \boldsymbol{v} + \frac{\gamma}{\gamma - 1} Re^{-1} Pr^{-1} \nabla T + S^{-1}\left(\frac{1}{2}\nabla(\boldsymbol{B} \cdot \boldsymbol{B}) - B(\nabla \boldsymbol{B})^{\mathrm{T}}\right) \end{bmatrix} \tag{10}$$

with

$$\boldsymbol{\tau} = \nabla \boldsymbol{v} + (\nabla \boldsymbol{v})^T - \tfrac{2}{3} \nabla \cdot \boldsymbol{v} I,$$

where we have assumed a spatially homogeneous viscosity μ, resistivity η, and thermal conductivity κ. The non-dimensional parameters governing the solution of this system are the Reynolds number $Re = \rho_0 v_A l_0 / \mu$, the Lundquist number $S = \mu_0 v_A l_0 / \eta$ and the Prandtl number $Pr = c_p \mu / \kappa$ for a characteristic density ρ_0, a characteristic length scale l_0, and the Alfven velocity v_A (as usual the permeability of free space is denoted by μ_0, $\kappa = 5/3$, and c_p is the specific heat of the fluid).

Note that in all the simulations we use dimensionless units. Thus, we can rewrite (10) more conveniently in terms of the (dimensionless) viscosity $\mu = Re^{-1}$, the (dimensionless) resistivity $\eta = S^{-1}$, and the (dimensionless) thermal conductivity $\kappa = Pr^{-1}$. Then, the form of the resistive MHD

equations used for the spatial discretization is

$$\frac{\partial U}{\partial t} + \nabla \cdot F(U) = \nabla \cdot F_d(U)$$

with $F(U)$ given by equation (9) and where $F_d(U)$ is given by

$$F_d(U) = \begin{bmatrix} 0 \\ \mu \boldsymbol{\tau} \\ \eta \left(\nabla B - (\nabla B)^{\mathrm{T}} \right) \\ \mu \boldsymbol{\tau} \cdot \boldsymbol{v} + \frac{\gamma \mu \kappa}{\gamma - 1} \nabla T + \eta \left(\frac{1}{2} \nabla (\boldsymbol{B} \cdot \boldsymbol{B}) - B (\nabla \boldsymbol{B})^{\mathrm{T}} \right) \end{bmatrix}.$$

For the numerical simulations conducted in this paper, we employ the MHD code developed in [11]. This C++ code is based on the Fortran code developed in [38] which has been used to conduct plasma physics simulations (see e.g. [38], [39], and [37]) as well as to construct more efficient preconditioners in the context of implicit time integrators (see [39]).

Our implementation assumes a finite difference or finite volume method where a single value is stored in each cell. Thus, in each cell we store the value of the density ρ, the fluid velocity $\boldsymbol{v}$, the magnetic field $\boldsymbol{B}$, the energy e (but not the pressure p). For the numerical simulations conducted in this paper, we have implemented the 2.5-dimensional case (that is, the state variables only depend on the x- and y-direction but both the velocity $\boldsymbol{v}(x, y)$ and the magnetic field $\boldsymbol{B}(x, y)$ are three-dimensional vectors) using a classic centered stencil for the divergence; that is, for each vector G, corresponding to the flux of a given (scalar) state variable, we compute the following approximation of the divergence

$$\nabla \cdot G(U) \approx \frac{G(U_{i+1,j}) - G(U_{i-1,j})}{2h} + \frac{G(U_{i,j+1}) - G(U_{i,j-1})}{2h},$$

where i and j are the cell indices in the x- and y-directions, respectively and h is the cell size. To evaluate the spatial derivatives present in the diffusion vector F_d, we also employ the classic centered stencil. Thus, the code has order two accuracy in space.

It is well known that for MHD problems preserving the divergence free property of the magnetic field is important. The space discretization used in our implementation ensures that if the solenoidal property $\nabla \cdot \boldsymbol{B} = 0$ is satisfied for the initial value, then this is true for all later times. In [39] it is shown that this property also holds true if we couple the space discretiza-

tion with an implicit method that is based on a matrix-free inexact Newton–Krylov algorithm. In addition, the numerical simulations conducted suggest that the divergence free condition is preserved if an exponential integrator is used for the time discretization.

4. The reconnection problem

The examples in this and the next section are drawn from [38] and [39], respectively. We start with a reconnection problem for which the initial value of the magnetic field is given by

$$B_0(x,y,z) = \begin{bmatrix} \tanh(2y) - \psi_0 k_y \cos(k_x x)\sin(k_y y) \\ \psi_0 k_x \sin(k_x x)\cos(k_y y) \\ 0 \end{bmatrix},$$

where as in [38] we have chosen $k_x = \pi/x_r$, $k_y = \pi/(2y_r)$, $\psi_0 = 0.1$, and the computational domain is given by $[-x_r, x_r] \times [-y_r, y_r]$ for $x_r = 12.8$ and $y_r = 6.3$. This implies that the magnetic field reverses direction from pointing along $\boldsymbol{e}_x$ to pointing along $-\boldsymbol{e}_x$ abruptly at $y = 0$. Furthermore, we impose a density that is given by

$$\rho = 1.2 - \tanh^2(2y)$$

and a pressure that is proportional to the density; to be more precise $p = 0.5\rho$ (from which the energy is determined). A vanishing velocity in both space directions is prescribed.

The time evolution for the current is shown in Figure 1. In all the simulations conducted the relative tolerance, the absolute tolerance, and the tolerance for the Krylov iteration are equal. We have investigated the effect of varying the tolerance for the Krylov iteration and found no significant difference in accuracy or run time as long as we choose a tolerance for the Krylov iteration that is at least as small as the tolerance for the time integration.

First, we consider the configuration that is investigated in [38]; that is, we employ 256 grid points in the x-direction and 128 grid points in the y-direction. The results obtained are shown in Figure 2. We observe that the performance of the EPIRK5P1 method is superior to CVODE if the desired accuracy is equal or less than 10^{-4}. In the low precision regime that is often of interest in applications, EPIRK5P1 outperforms CVODE by almost a factor of 2. However, for more stringent tolerances the CVODE

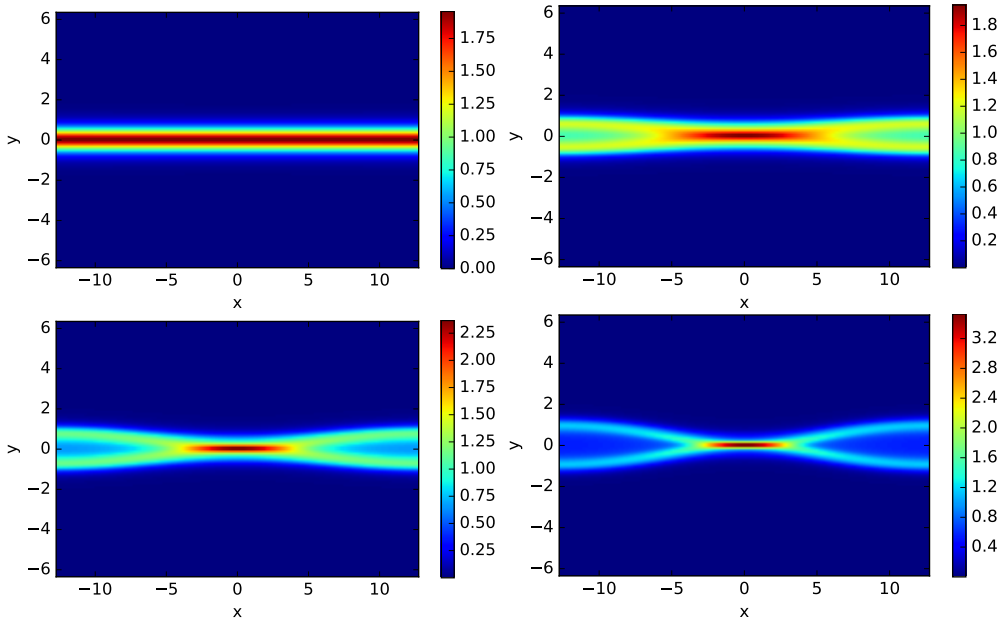

Figure 1: The current $j = \nabla \times B$ for the reconnection problem is shown at times $t = 0$ (top-left), $t = 100$ (top-right), $t = 150$ (bottom-left), and $t = 200$ (bottom-right). We employ 256 grid points in the x-direction and 128 grid points in the y direction. The dimensionless parameters are chosen as follows: $\mu = 10^{-2}$, $\eta = 10^{-3}$, and $\kappa = 10^{-2}$. The absolute and relative tolerance for EPIRK5P1 is set to 10^{-6}.

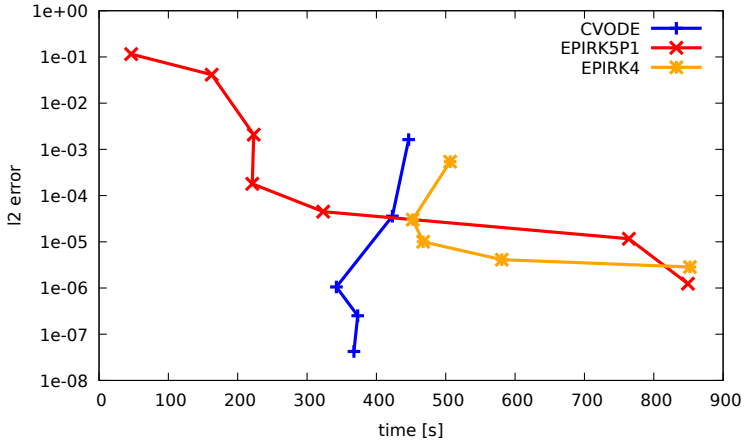

Figure 2: The maximum error of the components of the numerical solution (measured in the l^2 norm) is shown as a function of the run time. The reference solution for the CVODE/EpiRK is computed using the EpiRK/CVODE method with a tolerance of 10^{-14}. The simulation is conducted up to a final time of $T = 100$. We have used 256 grid points in the x-direction and 128 grid points in the y-direction (this corresponds to the configuration investigated in [38]). The viscosity is given by $\mu = 5 \cdot 10^{-2}$, the resistivity by $\eta = 5 \cdot 10^{-3}$, and the thermal conductivity by $\kappa = 4 \cdot 10^{-2}$. CVODE tolerance: 10^{-4}, 10^{-5}, 10^{-7}, 10^{-8}, 10^{-9}. EPIRK5P1 tolerance: 10^{-1}, $2 \cdot 10^{-2}$, 10^{-2}, 10^{-3}, 10^{-5}, 10^{-7}, 10^{-9}. EPIRK4 step size: 2, 1, $4 \cdot 10^{-1}$, $2 \cdot 10^{-1}$, 10^{-1}.

implementation is significantly more efficient. This is due to the fact that the run time for the CVODE method is almost independent of the accuracy achieved. What is perhaps even more interesting is that for a specified tolerance of 10^{-4} or above the CVODE method does not converge (in the situation described a final step size on the order of 10^{-16} is reported by the application). Let us also note that the constant time step method EPIRK4 is not competitive for this problem.

Next, we investigate the effect of increasing the number of grid points. The results obtained with 512 grid points in the x-direction and 256 grid points in the y-direction are shown Figure 3. In this case almost identical conclusions can be drawn. Once again the performance of EPIRK5P1 is superior to CVODE library if the desired accuracy is equal or less than 10^{-4}. Furthermore, the constant step size method EPIRK4 is not competitive except for very stringent tolerance requirements.

Up to this point we have only considered relatively large values for the

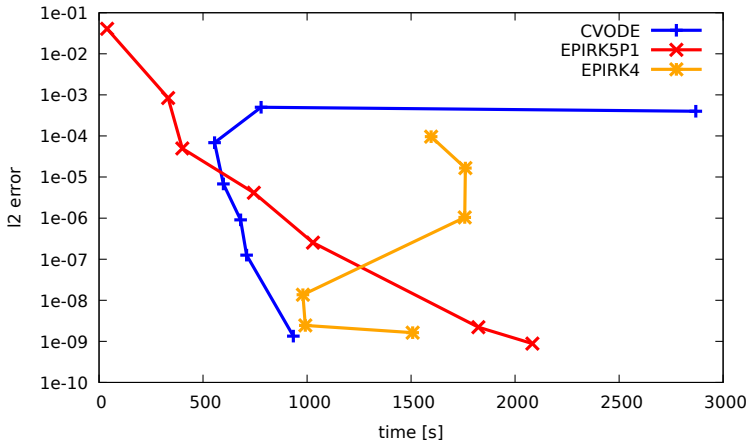

Figure 3: The maximum error of the components of the numerical solution (measured in the l^2 norm) is shown as a function of the run time. The reference solution for the CVODE/EpiRK is computed using the EpiRK/CVODE method with a tolerance of 10^{-11}. The simulation is conducted up to a final time of $T = 20$. We have used 512 grid points in the x-direction and 256 grid points in the y-direction. The viscosity is given by $\mu = 5 \cdot 10^{-2}$, the resistivity by $\eta = 5 \cdot 10^{-3}$, and the thermal conductivity by $\kappa = 4 \cdot 10^{-2}$. CVODE tolerance: 10^{-3}, 10^{-4}, ..., 10^{-9}. EPIRK5P1 tolerance: 10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}, 10^{-8}, 10^{-9}. EPIRK4 step size: 4, 2, 1, $4 \cdot 10^{-1}$, $2 \cdot 10^{-1}$, 10^{-1}.

viscosity, resistivity, and the thermal conductivity. Thus, we proceed by decreasing these dimensionless quantities by a factor of five. The corresponding results are shown in Figure 4. In this case the EPIRK5P1 implementation outperforms the CVODE implementation for all the tolerances studied here (although the difference in performance between the two methods in the low precision regime is smaller than was the case for the previous two examples). For accuracies between 10^{-6} and 10^{-8} the constant time step method EPIRK4 manages to outperform both the EPIRK5P1 as well as CVODE implementation.

We observe that for larger Reynolds and Lundquist numbers the EpiRK method shows better or identical performance for the range of tolerances studied here. In addition, the EpiRK method shows a clear advantage over the CVODE implementation for tolerances larger than 10^{-5}.

In summary, it can be said that for the reconnection problem the EPIRK5P1 implementation shows superior performance in the large and medium tolerance range that is of interest for the majority of practical applications.

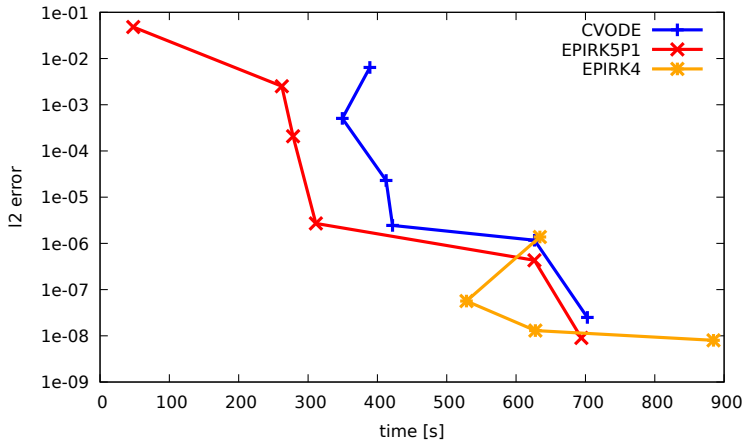

Figure 4: The maximum error of the components of the numerical solution (measured in the l^2 norm) is shown as a function of the run time. The reference solution for the CVODE/EpiRK is computed using the EpiRK/CVODE method with a tolerance of 10^{-14}. The simulation is conducted up to a final time $T = 100$. We have used 256 grid points in the x-direction and 128 grid points in the y-direction. The viscosity is given by $\mu = 10^{-2}$, the resistivity by $\eta = 10^{-3}$, and the thermal conductivity by $\kappa = 10^{-2}$. CVODE tolerance: 10^{-4}, 10^{-5}, ..., 10^{-9}. EPIRK5P1 tolerance: 10^{-1}, 10^{-2}, 10^{-3}, 10^{-5}, 10^{-7}, 10^{-8}. EPIRK4 step size: 2, 1, $2 \cdot 10^{-1}$, 10^{-1}. Note that CVODE does not converge for tolerances above 10^{-4}.

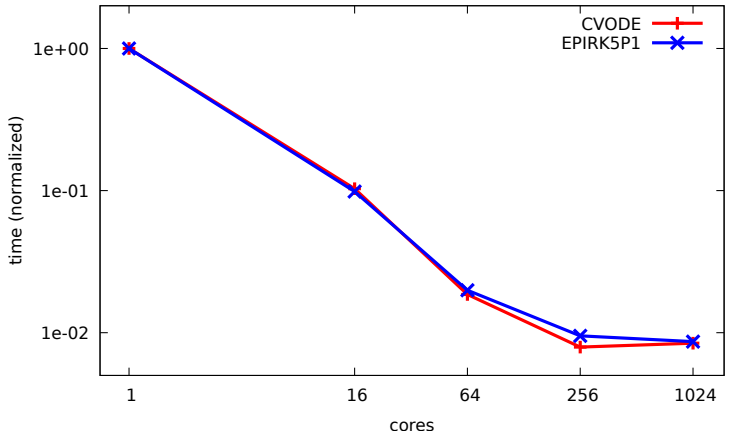

Figure 5: The run time (normalized to the single core run time for both the CVODE and EpiRK implementation) is shown as a function of the number of cores. The simulation is conducted up to a final time $T = 1$. We have used 512 grid points in both the x and y-direction. The viscosity is given by $\mu = 5 \cdot 10^{-2}$, the resisitivity by $\eta = 5 \cdot 10^{-3}$, and the thermal conductivity by $\kappa = 4 \cdot 10^{-2}$. These performance measurements were done on the VSC-2 (Vienna Scientific Cluster 2; see `http://vsc.ac.at/systems/vsc-2/`) which uses two 8 core AMD Opteron Magny Cours 6132HE and 32 GB of DDR3 memory per node. The nodes are connected by Infiniband QDR and the Intel MPI library (version 4.1) is employed.

However, for some of the examples considered here the CVODE implementation has a significant advantage for more stringent tolerance requirements.

To conclude this section let us investigate the strong scaling (i.e. the problem size is kept fixed while the number of cores is increased) for both our time integration routines and the CVODE implementation. This allows us to ascertain how well the current implementation is parallelized. In the ideal case we would observe a linear decrease in the run time as a function of the number of cores.

The strong scaling results on the VSC-2 are shown in Figure 5 and are almost identical for both implementations. We conclude that the algorithms scale well up to approximately 256 cores. No further gain in performance can be observed for 1024 or more cores. This is to be expected as in the case of strong scaling the amount of work available to a single core eventually decreases sufficiently such that communication between he different MPI processes limits the performance of the application. At this point no further increase in performance can be expected.

parameter	value
ϵ_x, ϵ_y	0.1
v_0	0.5
ω_x, ω_y	2
pressure p	0.25
B_z	10
B_x	0.1

Table 2: Parameters for the initial value of the Kelvin-Helmholtz instability.

5. The Kelvin-Helmholtz instability

As a second example we consider the Kelvin-Helmholtz (KH) instability. The KH instability is triggered by superimposing the following perturbation

$$\epsilon_x \cos\left(\tfrac{2\pi\omega_x}{L_x}x\right) + \epsilon_y \sin\left(\frac{\pi(2\omega_y - 1)}{L_y}y\right)$$

in the x-direction on the velocity field

$$\boldsymbol{v} = \begin{bmatrix} v_0 \tanh(\frac{y}{\lambda}) \\ 0 \\ 0 \end{bmatrix}.$$

The density is initialized to unity and a uniform pressure is chosen. The magnetic field is initialized to be uniform in both the x- and z-direction with strength B_x and B_z, respectively, and is assumed to vanish in the y-direction. All the parameters used to determine the numerical value of the initial value are listed in Table 2.

The time evolution of the numerical solution is shown in Figure 6. In all the simulations conducted the relative tolerance, the absolute tolerance, and the tolerance for the Krylov iteration are equal. We have investigated the effect of varying the tolerance for the Krylov iteration and found no significant difference in accuracy or run time as long as we choose a tolerance for the Krylov iteration that is at least as small as the tolerance for the time integration.

First, we perform a numerical simulation of the Kelvin–Helmholtz instability for 128 grid points in both space directions. The results obtained are shown in Figure 7. Both exponential integrators considered here (EPIRK5P1 and EPIRK4) outperform the CVODE implementation by a

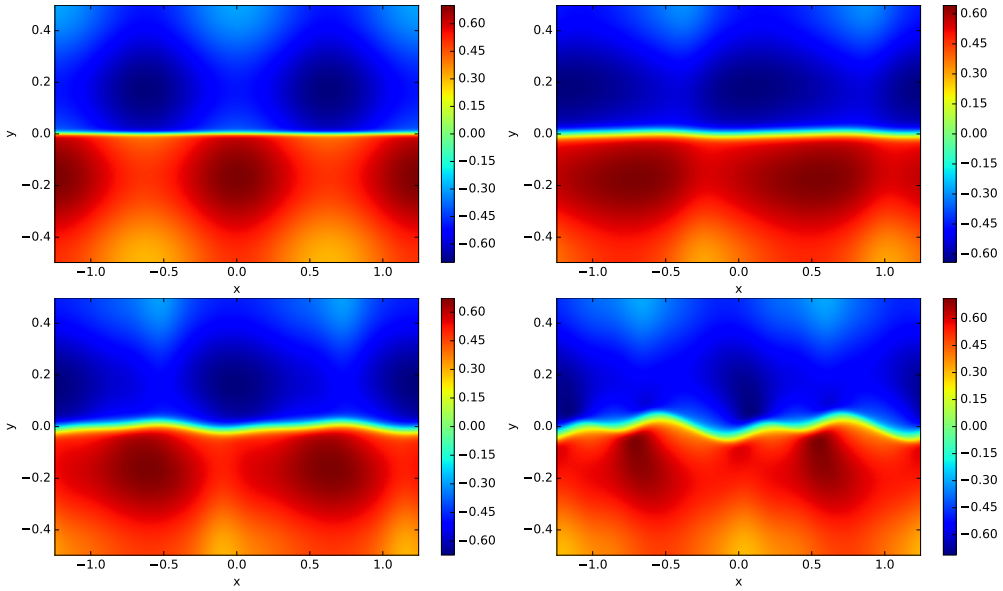

Figure 6: For the Kelvin-Helmholtz instability the velocity in the x-direction is shown at time $t = 0$ (top-left), $t = 1.5$ (top-right), $t = 2$ (bottom-left), and at $t = 2.5$ (bottom-right). In both space directions 256 grid points are employed. The dimensionless parameters are chosen as follows: $\mu = 10^{-4}$, $\eta = 10^{-4}$, and $\kappa = 10^{-4}$. The absolute and relative tolerance for EPIRK5P1 is set to 10^{-6}.

significant margin (approximately a factor of two). Furthermore, CVODE is very sensitive to even small perturbations in the chosen tolerance, while the exponential integrators are much more robust in this regard. Thus, a user who would naively (i.e. without measuring the run time for a range of tolerances) choose a specific tolerance would increase the run time of the application by significantly more than the factor of two stated earlier. What is also interesting to see is that the constant step size method EPIRK4 is equal in performance or outperforms EPIRK5P1 for all tolerances studied here (although the difference between these two methods is at most 30%).

Now, let us increase the number of grid points to 256 (in both space directions). The corresponding results are shown in Figure 8. As before, we observe that both exponential integrators significantly outperform the CVODE implementation. In addition, the constant step size method EPIRK4 outperforms the EPIRK5P1 implementation by up to 30%. This is also consistent with the results on the coarse grid considered before. Let us emphasize here again that the major advantage of EPIRK4 is that it uses fewer Krylov projections (two vs. three for EPIRK5P1) per time step. However, at present this advantage over EPIRK5P1 is limited to constant time step implementations. A third projection is needed to construct automatic time stepping control schemes for EPIRK4. The variable time stepping version of the integrators is more efficient for EPIRK5P1 since it is a higher order method. For the Kelvin–Helmholtz instability, considered in this section, these differences work in favor of the EPIRK4 method while for the reconnection problem, considered in the previous section, the EPIRK5P1 scheme is superior (in the case of low and medium tolerances).

In summary, both exponential integrators considered in this paper show significantly improved performance compared to CVODE for the Kelvin–Helmholtz instability. Let us also mention that for both the exponential integrators as well as CVODE it is not entirely straight forward to choose a good tolerance. This is due to the fact that if the tolerance is chosen too high, the run time can actually increase. We emphasize, however, that this behavior is significantly more serious for the CVODE implementation.

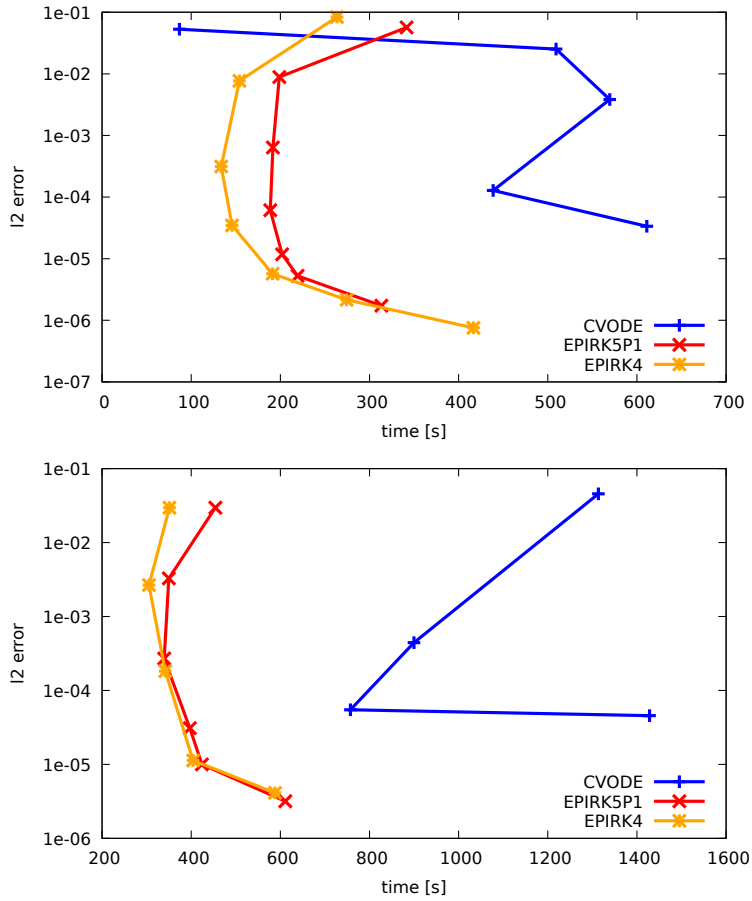

Figure 7: The maximum error of the components of the numerical solution (measured in the l^2 norm) is shown as a function of the run time. The reference solution for the CVODE/EpiRK is computed using the EpiRK/CVODE method with a tolerance of 10^{-11}. The simulation is conducted up to a final time $T = 1$ (top) and $T = 2$ (bottom) and 128 grid points are used in both the x- and y-direction. The (dimensionless) viscosity is given by $\mu = 10^{-4}$, the resistivity by $\eta = 10^{-4}$, and the thermal conductivity by $\kappa = 10^{-4}$. CVODE tolerance ($T = 1$): 10^{-3}, 10^{-6}, 10^{-7}, 10^{-8}, 10^{-9}. CVODE tolerance ($T = 2$): 10^{-6}, 10^{-7}, 10^{-8}, 10^{-9}. EPIRK5P1 tolerance ($T = 1$): 10^{-3}, 10^{-4}, ..., 10^{-9}. EPIRK5P1 tolerance ($T = 2$): 10^{-4}, 10^{-5}, ..., 10^{-9}. EPIRK4 step size ($T = 1$): 10^{-1}, $4 \cdot 10^{-2}$, $2 \cdot 10^{-2}$, 10^{-2}, $4 \cdot 10^{-3}$, $2 \cdot 10^{-3}$, 10^{-3}. EPIRK4 step size ($T = 2$): $4 \cdot 10^{-2}$, $2 \cdot 10^{-2}$, 10^{-2}, $4 \cdot 10^{-3}$, $2 \cdot 10^{-3}$.

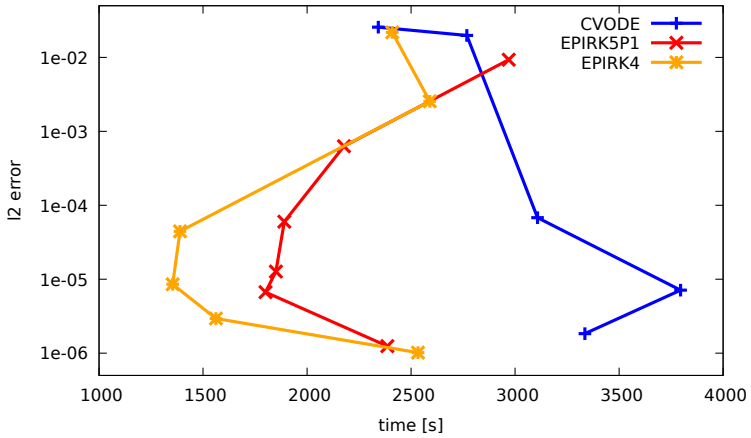

Figure 8: The maximum error of the components of the numerical solution (measured in the l^2 norm) is shown as a function of the run time. The reference solution for the CVODE/EpiRK is computed using the EpiRK/CVODE method with a tolerance of 10^{-11}. The simulation is conducted up to a final time $T = 1$ and 256 grid points are used in both the x- and y-direction. The (dimensionless) viscosity is given by $\mu = 10^{-4}$, the resistivity by $\eta = 10^{-4}$, and the thermal conductivity by $\kappa = 10^{-4}$. CVODE tolerance: 10^{-5}, 10^{-6}, ..., 10^{-9}. EPIRK5P1 tolerance: 10^{-4}, 10^{-5}, ..., 10^{-9}. EPIRK4 step size: $4 \cdot 10^{-2}$, $2 \cdot 10^{-2}$, 10^{-2}, $4 \cdot 10^{-3}$, $2 \cdot 10^{-3}$, 10^{-3}.

6. Implementation

We have chosen to base our implementation on the 2.5 dimensional[1] MHD code developed by Daniel R. Reynolds et al.. A large number of numerical simulations have been conducted (see e.g. [38] and [39]) and the performance of preconditioners for Newton-Krylov based implicit methods has been investigated (see e.g. [39]). Furthermore, the extension of the code to three dimensional problems as well as to non-square geometries (such as a Tokamak geometry) have been investigated in [37].

Our code is implemented in C++ and is designed to easily accommodate different time integrators, space discretization schemes, as well as different MHD models. The design considerations and a detailed description of the computer code can be found in [11].

The second component of our code is the EPIC library (see [28]) which implements the fifth order EpiRK5P1 with an adaptive time step control

[1] All quantities depend on only two spatial variables; however, the direction of the velocity as well as the magnetic fields are three dimensional vectors.

(*EpiRK5P1VerticalVar* function in EPIC library), the fourth order EpiRK4 method (*EpiRK4MixedConst* function in EPIC library) with constant time steps and the adaptive Krylov algorithm to evaluate the φ_k functions as described in [34]. This library is written in the C++ programming language with MPI, and uses routines from the BLAS and LAPACK libraries. In the experiments, the BLAS and LAPACK implementations were supplied by the Intel Math Kernel Library (Intel MKL).

One detail warrants further discussion: in the original Fortran program the CVODE interface is used to approximate the Jacobian by a simple forward difference stencil. In our implementation, however, we provide a custom function to compute an approximation to the Jacobian. This is necessary for the EPIC library, which before was only used in the context of problems where an exact Jacobian is available (as in [28]). Computing an exact Jacobian for the significantly more complex MHD problem considered in this paper is infeasible. Therefore, we have to approximate $J(a)v$, i.e. the application of the Jacobian J at position a to a vector v, where both a and v depend on the specific numerical algorithm as well as the initial values. We note here, that there is a qualitative difference in the norm of the vector v that depends on the numerical time integration scheme used. In the EpiRK the norm of v is close to unity whereas in the BDF method the value is often significantly below $\sqrt{\epsilon}$ (where we use $\epsilon \approx 10^{-16}$ to denote machine precision). Therefore, care is taken to scale vectors to the norm $\sqrt{\epsilon}$ only if the initial norm is above $\sqrt{\epsilon}$ in magnitude. Then the same implementation can be used for both methods and a difference in performance due to an internally optimized function to compute the Jacobian (as is provided by the CVODE library) is precluded.

In all of the computations conducted in this paper, we have used the infinity norm to scale v to its appropriate size. In general, we have found that the performance and accuracy of the computation is not significantly altered if the scaling is performed to some value that is reasonably close to $\sqrt{\epsilon}$.

7. Conclusion

Using several examples of MHD problems, we have shown that exponential integrators constitute a viable alternative to the more commonly employed BDF method (as implemented in the CVODE library). In almost all instances equal or superior performance has been observed for the

adaptive and variable time stepping fifth order EpiRK method for low to medium accuracy requirements.

Note that while the first version of the CVODE package was released several decades ago and the package has been optimized and refined over extended period of time, the EPIC package is the first C++/MPI implementation of the exponential integrators. As indicated in [29] there are a number of additional optimizations possible both from an algorithmic and from a computer science perspective. For example, significant performance gains are expected from improvements in the adaptivity algorithms in the Krylov projections and the derivation of the time integrators specifically adapted to a given system of equations. In addition, the newly introduced implicit-exponential (IMEXP) methods [36] demonstrate that exponential integration can also be used as a component of the overall integration scheme and can be combined with the preconditioned implicit integrators as well as time splitting strategies. It is expected that further development will lead to significant performance improvements.

Furthermore, we note that the results presented here are somewhat different from the simpler models considered in [29]. In that case order of magnitude speedups are observed for the EpiRK method as compared to the CVODE implementation. Further investigations are needed to determine whether the degradation in comparative performance is due to the specifics of the PDE system, the details of the integrator or the particulars of the implementation. Also it is not clear to what degree the approximate computation of the Jacobian contributes to the observed differences (in all the problems investigated in [29] an analytical form of the Jacobian has been used).

Acknowledgements

We would like to take the opportunity to thank D. R. Reynolds for providing the code of the Fortran MHD solver and for the helpful discussion. Dr. Einkemmer was supported by the Marshall plan scholarship of the Austrian Marshall plan foundation (http://www.marshallplan.at/) and by the Fonds zur Förderung der Wissenschaften (FWF) – project id: P25346. Dr. Tokman was supported by a grant from the National Science Foundation, Computational Mathematics Program, under Grant No. 1115978.

The simulations were conducted using the Vienna Scientific Cluster (VSC).

References

[1] A. H. Al-Mohy and N. J. Higham, *Computing the action of the matrix exponential, with an application to exponential integrators*, SIAM J. Sci. Comput. **33** (2011), no. 2, 488–511.

[2] W.E. Arnoldi, *The principle of minimized iteration in the solution of the matrix eigenvalue problem*, Quart. Appl. Math. **9** (1951), 17–29.

[3] L. Bergamaschi, M. Caliari, and M. Vianello, *The relpm exponential integrator for fe discretizations of advection-diffusion equations.*, International Conference on Computational Science (Marian Bubak, G. Dick van Albada, Peter M. A. Sloot, and Jack Dongarra, eds.), Lecture Notes in Computer Science, vol. 3039, Springer, 2004, pp. 434–442.

[4] J. Breslau, N. Ferraro, and S. Jardin, *Some properties of the M3D-C1 form of the three-dimensional magnetohydrodynamics equations*, Phys. Plasmas **16** (2009), no. 9, 822–838.

[5] M. Caliari and A. Ostermann, *Implementation of exponential Rosenbrock-type integrators*, Appl. Numer. Math. **59** (2009), 568–581.

[6] M. Caliari, M. Vianello, and L. Bergamaschi, *Interpolating discrete advection–diffusion propagators at Leja sequences*, J. Comput. Appl. Math. **172** (2004), no. 1, 79 – 99.

[7] L. Chacón, *An optimal, parallel, fully implicit Newton–Krylov solver for three-dimensional viscoresistive magnetohydrodynamicsa)*, Phys. Plasmas **15** (2008), no. 5, 056103.

[8] S.D. Cohen and A.C. Hindmarsh, *CVODE, a stiff/nonstiff ODE solver in C*, Comput. Phys. **10** (1996), no. 2, 138–143.

[9] H. A. Van der Vorst, *An iterative solution method for solving $f(A) = b$ using Krylov subspace information obtained for the symmetric positive definite matrix A*, J. Comput. Appl. Math. **18** (1987), 249–263.

[10] F. Ebrahimi and R. Raman, *Plasmoids formation during simulations of coaxial helicity injection in the national spherical torus experiment*, Phys. Rev. Lett. **114** (2015), 205003.

[11] L. Einkemmer, *A modern resistive magnetohydrodynamics solver using C++ and the Boost library*, arXiv preprint, arXiv:1407.3189 (2015).

[12] L. Einkemmer and A. Ostermann, *Exponential integrators on graphic processing units*, High Performance Computing and Simulation HPCS, 2013 International Conference on (2013).

[13] M.R. Fahey and J. Candy, *GYRO: A 5-D Gyrokinetic-Maxwell solver*, Proceedings of the ACM/IEEE SC 2004 Conference (2008).

[14] E. Frieman and L. Chen, *Nonlinear gyrokinetic equations for low-frequency electromagnetic waves in general plasma equilibria*, Phys. Fluids **25** (1982).

[15] B. Fryxell, K. Olson, P. Ricker, F. X. Timmes, M. Zingale, D. Q. Lamb, P. MacNeice, R. Rosner, J. W. Truran, and H. Tufo, *FLASH: an adaptive mesh hydrodynamics code for modeling astrophysical thermonuclear flashes*, Astrophys. J. Suppl. S. **131** (2000), no. 1, 273.

[16] H. Goedbloed and S. Poedts, *Principles of magnetohydrodynamics with applications to laboratory and astrophysical plasmas.*, Cambridge University Press, 2004.

[17] A. C. Hindmarsh, P. N. Brown, K. E. Grant, S. L. Lee, R. Serban, D. E. Shumaker,

and C. S. Woodward, *SUNDIALS: Suite of nonlinear and differential/algebraic equation solvers*, ACM Trans. Math. Softw. **31** (2005), no. 3, 363–396.

[18] M. Hochbruck and C. Lubich, *Exponential integrators for quantum-classical molecular dynamics*, BIT Num. Math. **39** (1999), no. 4.

[19] M. Hochbruck and Ch. Lubich, *On Krylov subspace approximations to the matrix exponential operator*, SIAM J. Numer. Anal. **34** (1997), 1911–1925.

[20] M. Hochbruck and A. Ostermann, *Explicit exponential Runge–Kutta methods for semilinear parabolic problems*, SIAM J. Numer. Anal. **43** (2005), no. 3.

[21] M. Hochbruck and A. Ostermann, *Exponential integrators*, Acta Numerica **19** (2010), 209–286.

[22] S.C. Jardin, *Review of implicit methods for the magnetohydrodynamic description of magnetically confined plasmas*, J. Comput. Phys. **231** (2012), no. 3, 822 – 838, Special Issue: Computational Plasma Physics.

[23] A.K. Kassam and L.N. Trefethen, *Fourth-order time stepping for stiff PDEs*, SIAM J. Sci. Comput. **26** (2005), no. 4, 1214–1233.

[24] R. Keppens, G. Toth, M.A. Botchev, and A. Van Der Ploeg, *Implicit and semi-implicit schemes: Algorithms*, Int. J. Numer. Meth. Fl. **30** (1999), no. 3, 335–352.

[25] C. Klein and K. Roidot, *Fourth order time-stepping for Kadomtsev-Petviashvili and Davey-Stewartson equations*, SIAM J. Sci. Comput. **33** (2011), no. 6.

[26] T. Kuwabara, K. Shibata, T. Kudoh, and R. Matsumoto, *The acceleration mechanism of resistive magnetohydrodynamic jets launched from accretion disks*, Astrophys. J. **621** (2005), no. 2, 921.

[27] I. Lindemuth and J. Killeen, *Alternating direction implicit techniques for two-dimensional magnetohydrodynamic calculations*, J. Comput. Phys. **13** (1973), no. 2, 181 – 208.

[28] J. Loffeld and M. Tokman, *Comparative performance of exponential , implicit , and explicit integrators for stiff systems of ODEs*, J. Comput. Appl. Math. (2012), 45–67.

[29] ———, *Implementation of parallel adaptive-Krylov exponential solvers for stiff problems*, SIAM J. Sci. Comput. **36** (2014), no. 5, C591–C616.

[30] H. Lütjens and J. Luciani, *XTOR-2F: a fully implicit Newton-Krylov solver applied to nonlinear 3D extended MHD in tokamaks*, J. Comput. Phys. **229** (2010), no. 21, 8130–8143.

[31] A. Martinez, L. Bergamaschi, M. Caliari, and M. Vianello, *A massively parallel exponential integrator for advection-diffusion models*, J. Comput. Appl. Math. **231** (2009), no. 1.

[32] C.B. Moler and C.F. Van Loan, *Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later*, SIAM Rev. **45** (2003), no. 1, 3–49.

[33] D.R. Nicholson, *Introduction to plasma theory*, Cambridge University Press, 1983.

[34] J. Niesen and W.M. Wright, *Algorithm 919: a Krylov subspace algorithm for evaluating the φ-functions appearing in exponential integrators*, ACM T. Math. Software **38** (2012), no. 3.

[35] G. Rainwater and M. Tokman, *A new class of split exponential propagation iterative methods of Runge–Kutta type (sEPIRK) for semilinear systems of ODEs*, J. Comput. Phys. **269** (2014), 40 – 60.

[36] ———, *A new approach to constructing efficient stiffly accurate epirk methods*, (under review) (available at INSERT ARXIV LINK).

[37] D. R. Reynolds, R. Samtaney, and H. C. Tiedeman, *A fully implicit Newton–Krylov–Schwarz method for tokamak magnetohydrodynamics: Jacobian construction and preconditioner formulation*, Comput. Sci. Discov. **5** (2012), no. 1, 014003.

[38] D. R. Reynolds, R. Samtaney, and C. S. Woodward, *A fully implicit numerical method for single-fluid resistive magnetohydrodynamics*, J. Comput. Phys. **219** (2006), no. 1, 144 – 162.

[39] D.R. Reynolds, R. Samtaney, and C.S. Woodward, *Operator-based preconditioning of stiff hyperbolic systems*, SIAM J. Sci. Comput. **32** (2010), no. 1, 150–170.

[40] Y. Saad, *Analysis of some Krylov subspace approximations to the matrix exponential operator*, SIAM J. Numer. Anal. **29** (1992), 209–228.

[41] ———, *Iterative methods for sparse linear systems*, PWS Publishing Company, 1996.

[42] D. Schnack and J. Killeen, *Nonlinear, two-dimensional magnetohydrodynamic calculations*, J. Comput. Phys. **35** (1980), no. 1, 110 – 145.

[43] R.B. Sidje, *Expokit: A software package for computing matrix exponentials*, ACM Trans. Math. Softw. **24** (1998), 130–156.

[44] C.R. Sovinec, A.H. Glasser, T.A. Gianakon, D.C. Barnes, R.A. Nebel, S.E. Kruger, S.J. Plimpton, A. Tarditi, M.S. Chu, and the NIMROD Team, *Nonlinear magnetohydrodynamics with high-order finite elements*, J. Comp. Phys. **195** (2004), 355.

[45] C.R. Sovinec and J.R. King, *Analysis of a mixed semi-implicit/implicit algorithm for low-frequency two-fluid plasma modeling*, J. Comput. Phys. **229** (2010), no. 16, 5803 – 5819.

[46] M. Tokman, *Magnetohydrodynamic modeling of solar magnetic arcades using exponential propagation methods*, Ph.D. thesis, California Institute of Technology, 2001.

[47] M. Tokman, *Efficient integration of large stiff systems of ODEs with exponential propagation iterative (EPI) methods*, J. Comput. Phys. **213** (2006), no. 2, 748–776.

[48] ———, *A new class of exponential propagation iterative methods of Runge–Kutta type (EPIRK)*, J. Comput. Phys. **230** (2011), no. 24, 8762–8778.

[49] ———, *Four classes of exponential EPIRK integrators*, Oberwolfach Reports **14** (2014), 855–858.

[50] M. Tokman and P. M. Bellan, *Three-dimensional model of the structure and evolution of coronal mass ejections*, Astrophys. J. **567** (2002), no. 2, 1202.

[51] M. Tokman, J. Loffeld, and P. Tranquilli, *New adaptive exponential propagation iterative methods of Runge-Kutta type (EPIRK)*, SIAM J. Sci. Comput. **34** (2012), no. 5, A2650–A2669.

[52] ———, *New adaptive exponential propagation iterative methods of Runge-Kutta type (EPIRK)*, SIAM J. Sci. Comput. **34** (2012), no. 5, A2650–A2669.

[53] C. Zanni, A. Ferrari, R. Rosner, G. Bodo, and S. Massaglia, *MHD simulations of jet acceleration from Keplerian accretion disks - The effects of disk resistivity*, Astron. Astrophys. **469** (2007), no. 3, 811–828.